

Do Your Friends Make You Smarter?: An Analysis of Social Strategies in Online Information Seeking

Brynn M. Evans

University of California, San Diego
9500 Gilman Drive
La Jolla, CA 92093
brynn.evans@gmail.com

Sanjay Kairam

Palo Alto Research Center, Inc.
3333 Coyote Hill Road
Palo Alto, CA 94116
sanjay.kairam@gmail.com

Peter Pirolli

Palo Alto Research Center, Inc.
3333 Coyote Hill Road
Palo Alto, CA 94116
pirolli@parc.com

ABSTRACT

Prior research in the social search space has focused on the informational benefits of collaborating with others during web and workplace information seeking. However, social interactions, especially during complex tasks, can have cognitive benefits as well. Our goal in this paper is to document the methods and outcomes of using social resources to help with exploratory search tasks. We used a talk-aloud protocol and video capture to explore the actions of eight subjects as they completed two “Google-hard” search tasks. Task questions were alternated between a Social and Non-Social Condition. The Social Condition restricted participants to use only social resources—search engines were not allowed. The Non-Social Condition permitted normal web-based information sources, but restricted the use of social tools.

We describe the social tactics our participants used in their search process. Asking questions on social networking sites and targeting friends one-on-one both resulted in increased information processing but during different phases of the *question-answering* process. Participants received more responses via social networking sites but more thorough answers in private channels (one-on-one). We discuss the possibility that the technological and cultural affordances of different social-informational media may provide complementary cognitive benefits to searchers.

Our work suggests that online social tools could be better integrated with each other and with existing search facilities. We conclude with a discussion of our findings and implications for the design of social search tools.

Author Keywords

Social search, social information seeking, question-answering, information processing, cognitive task analysis, verbal protocols.

INTRODUCTION & RELATED WORK

As search engines and search algorithms have grown more sophisticated in recent years, people have become increasingly reliant on these technologies for seeking information. At the same time, the library science (Taylor 1968, Shepherd 1983, Fox et al. 1993, Twidale et al. 1997; Ehrlich & Cash 1994), educational (Rogoff 1990, Pea 1993, Hatch & Gardner 1993, Edelson, Pea & Gomez 1996, Karasavvidis 2002), and organizational learning literatures (Granovetter 1973, Allen 1977, Burt 1992, Cross & Sproull 2004) have documented how people also serve as valuable information resources. Cross and Sproull (2004) noted that in the workplace, 38 of the 40 managers interviewed “relied heavily on other people when seeking information, despite a high-quality computer-based knowledge repository” (p. 6). In particular, weak ties (Granovetter 1973), who tend to bridge across organizational and geographic boundaries, are known to supply new information, diverse ideas, and novel perspectives (Granovetter 1982, Burt 1992).

While online search has previously been thought of as a primarily solitary activity (Catledge & Pitkow 1995, Cockburn & Jones 1996), web searching and browsing can be a collaborative process (Twidale et al. 1997, Morris 2008, Evans & Chi 2008, Chi 2009). Both Morris (2008) and Evans and Chi (2008) discuss the prevalence and importance of social interactions during web search. As many as half of searchers may turn to others for advice, brainstorming opportunities, or search tips prior to embarking on a search (Evans & Chi 2008); social outlets are important following search, as well, for sharing information and seeking feedback. Additionally, it has been shown that peer support provides the greatest benefit when users are performing *informational* searches (Evans & Chi 2008)—especially when exploring ill-structured problem spaces with open-ended goals (Marchionini 2006). This type of exploratory searching can be hard to support from a system’s perspective due to the occasional gulf between users’ concepts and keywords and the jargon of the problem domain. Human-to-human communication has the potential to address this “vocabulary problem” (Furnas et al. 1987), highlighting the fact that social inputs may serve as cognitive aids during search.

Today’s web search tools provide relatively poor support for explicit social interactions during search tasks. SearchTogether (Morris & Horvitz 2007), CoSearch (Amershi & Morris 2008), and Cerchiamo (Pickens et al. 2008) are notable exceptions, although they were intentionally designed to support joint search tasks where two or more searchers have a common goal.

Yet information seekers with personal (non-collaborative) search goals may still engage in momentary collaborative efforts by emailing or calling friends for help. Can we also design integrated social search tools to support this class of user? To answer this question, we must first understand the specific benefits of social activities, social inputs, and collaboration in web-based information seeking.

We can draw some assumptions about these benefits by turning to the organizational learning literature, where Borgatti and Cross have proposed a formal model of information seeking in social contexts (2003). Their work suggests that people will seek the help of others who have domain expertise, who are accessible (or available to provide information in a timely manner), and who are willing to engage in problem solving (Cross & Sproull 2004). Cross et al. (2001) further suggest that social relationships will provide informational (e.g., solutions, facts, answers) and cognitive benefits (e.g., problem reformulation, problem framing) to searchers. While we might expect that this model applies in digital environments, neither the constraints nor the benefits of operating in online social contexts have been observed directly. However, we expect that peers and colleagues will provide online searchers with both informational and cognitive (reformulation, framing) benefits.

Our goal with this work is primarily exploratory. We hope to document the ways that searchers exploit their social environments during exploratory search tasks—but by observing social events that extend beyond web-based interactions. We are especially interested in users’ strategies for tapping their social networks, their social reasoning processes, and how *socio-cognitive* benefits apply during online sensemaking and problem solving. In what specific ways do social resources augment the search process? To look at this, we selectively recruited users with large and active social networks. We instructed them during one search task to use only social sources (Google, Yahoo!, MSN were not allowed). From this, we observed that *questioning* a social network and separately engaging in one-on-one *answering* sessions may help users think through problems. Long-term, we hope that this and future work can offer suggestions for the design of search tools that incorporate social support for users in situationally-appropriate ways.

METHOD

We used a talk-aloud protocol and video capture techniques to explore how eight users performed two search tasks related to U.S. energy policy. Below we describe our study design and data analysis procedures.

Subject Selection

Our call for participation went out formally on Craigslist and Stanford University (SU) Post, and informally through Twitter. All candidates (N=30) filled out pre-test questionnaires on their backgrounds and interests in energy policy, as well as their typical use of computers, the Internet, social networking, and web search. Because our goal was to discover effective social strategies for searching, we selected only eight individuals from this pool who were both highly social and skilled at web search—falling in the upper quartile of reported social networking activities and search behaviors from the set of 30 candidates.

Participants

Subjects (5 female, 3 male) came from a range of backgrounds and experiences (Table 1). All were between 21–40 years old, well-educated, and fairly liberal in their political leanings. Although the search tasks we selected could have been politically polarizing, politics did not come up as a major issue or impediment to reasoning during the study.

ID	Sex	Age	Education	Occupation
ss01	F	31–40	Bachelor’s	Program Manager
ss02	M	31–40	Master’s	Accountant
ss03	M	21–24	Master’s	Student
ss04	F	25–30	Bachelor’s	Finance
ss05	F	25–30	Bachelor’s	Information Technologist
ss06	M	31–40	Master’s	Consultant
ss07	F	25–30	Bachelor’s	Manager
ss08	F	31–40	Doctorate	Writer

Table 1. Participant Demographics (Note: we only asked for subject ages in these groupings; we do not have exact ages.)

Social network characteristics

While subjects were selected for being highly social, the nature of their social behaviors varied widely. Subjects' primary social networks were Twitter (N=5), LinkedIn (N=2), and Facebook (N=2) (one subject nominated both Twitter and LinkedIn equally); and self-reported network size ranged from 55 to 1000 (mean: 416; standard deviation: 323).

Social network diversity

In addition to network size, we estimated network diversity (or access to knowledge resources) through a *position generator* (Ganzeboom et al. 1992, Van der Gaag et al. 2008). This instrument lists 30 occupations of varying prestige, and for each occupation has users report whether it is held by a member of their social network (family, friend, or acquaintance). (Prestige scores were previously standardized to ISEI socioeconomic index measures (Ganzeboom et al. 1992)). We modified our instrument slightly from the one reported in Van der Gaag et al. (2008) and Marlow (2005) to include “environmental consultant” and “oil/gas industry worker,” as these occupations were relevant to the domain of energy policy. A sum of the ISEI occupational prestige scores in a user's network became her diversity measure (range: 341–1295; mean: 920; standard deviation: 283). This shows that, as with network size, there was a wide range of knowledge resources in users' networks. This measure was not used in the subject selection process, but was included in subsequent correlational analyses.

Task Questions

Our specific task questions were chosen for several reasons. We wanted to observe the search process for *exploratory* queries, or ones that involve incremental learning and investigating (Marchionini 2006). Thus, we required questions that were unfamiliar and also “Google-hard.” We determined that a question was “Google-hard” if, even after several attempts in major search engines, very few result listings had high information scent. We also wanted questions that had no single correct answer, but instead required a good deal of information foraging (learning) and subsequent sensemaking (synthesis)—a nice complement to their “Google-hard” nature.

Finally, we wanted to preserve a sense of real-life relevancy. Since U.S. gas prices were rapidly increasing at the time of the study (mid-2008), we thought that issues related to energy policies may be suitable task questions. These are not meant to be canonical *social search* queries—they were selected as “Google-hard” questions within the domain of energy policy—but consequently, we believed they would generate interesting (social) explorations.

55MPH: If we lowered the U.S. national speed limit to 55 miles per hour (MPH) (89 km/hour), how many fewer barrels of oil would the U.S. consume every year?

Pyrolysis: What role does pyrolytic oil (or pyrolysis) play in the debate over carbon emissions?

In fact, there was surprisingly little information about either of these topics available through Google (or in Wikipedia) at the time of our study. Thus, we told participants that we cared more about their exploration processes than whether they discovered the “right” answer.

Procedure

All participants engaged in two search blocks without time restrictions—first in the Social and then the Non-Social (NS) Condition. Task questions were counterbalanced between these conditions. After each search block, we interviewed subjects about their strategies and behaviors. Total session time (both blocks) varied between 37 and 120 minutes.

To best simulate subjects' natural working environments, we ran most users in their homes or workplaces; two preferred to come to our laboratory. Participants used their personal laptops, mobile phones, and regular suite of online tools for the session. Only one person additionally used pen-and-paper. All on-screen actions were recorded with screencast software (either Silverback [Mac] or Camtasia [Windows]). A Flip video camera or MacBook computer was instrumented to capture off-screen actions from the subject's desk and workspace. This track was used to corroborate events from on-screen actions and to view mobile phone interactions.

The Social Condition

In the Social Condition, participants were allowed to use only social resources to search for information. Friends, colleagues, social networking sites, question-answer sites, and blogs were permitted, but traditional search engines (Google, Yahoo!, MSN) and Wikipedia were not.¹ We asked that any recruited collaborators abide by the same restrictions as participants.

¹Prior to our study, we conducted a brief poll on Mechanical Turk (N=100) to see how people perceived “social” versus “non-social” resources, primarily to classify ambiguous sites like Wikipedia. Respondents were extremely consistent: For example, blogs were overwhelmingly “social,” while newspapers were not. Wikipedia was deemed to be “non-social” because it was perceived as an information database, not a social platform. Consequently, we restricted subjects from using Wikipedia in the Social Condition.

Although this was impossible to enforce, when subjects first began interacting with their partners, we requested that they communicate the Social Condition’s restrictions and, at the same time, obtain verbal consent from their collaborators to be remote participants in the study.

The Non- Social Condition

The Non-Social Condition was more similar to a traditional search session. It permitted the use of non-social resources (e.g., search engines, information databases)—including Wikipedia—and disallowed the use of the social resources mentioned above.

Talk-aloud protocol

We collected verbal protocols from participants during each search block per standard talk-aloud techniques (Ericsson & Simon 1980). We first provided instructions on the talk-aloud method and conducted a warm-up task; during the actual search blocks, we only interrupted subjects to remind them of the procedure if they became silent for an extended period.

Interview and retrospective probing

At the completion of each search block, we conducted a semi-structured interview with retrospective probing, consisting of a set of standard questions as well as questions specific to each participant’s behaviors. These so-called “verbal probes” served to elicit information about the people, technology, tools, or idiosyncratic behaviors we observed during the search tasks, but for which we did not interrupt. We asked subjects about their relationships with each person they contacted in the Social Condition; we discussed the history and typical usage of the tools they recruited; and we tried to understand the expected practices and cultural norms surrounding the technologies used.

Data Analysis

All videos were transcribed and manually reviewed. We initially coded for the types of activities users performed: whether they were describing their processes, having social interludes with friends, actively seeking information that addressed their queries, navigating interface elements, etc.; and whether their problem solving was conducted individually or in collaboration with others. This fine-grained demarcation of problem spaces revealed higher-level patterns that proved more worthwhile for our analyses [see *Behavioral Patterns* in the Results section].

Scoring Performance

Task performance was scored and included in analyses in multiple ways. For each task question, we identified thirteen “facts” central to the understanding of each question (for a selection of facts, see Table 2). We selected these based on the advice of field experts as well as the description of important findings from government and academic reports on oil savings and car speed (United States Government Accountability Office 2008) and pyrolysis (Lehmann 2007).

55MPH	Pyrolysis
Unknown to what extent drivers will follow speed limit	The pyrolysis process takes biomass or organic matter as inputs
Peak fuel efficiency in most cars between 30–60MPH	The pyrolysis process is carbon-neutral or carbon-negative
Different states have different speed limits	The bio-oil produced still requires refining before use
Most roads already have speed limits of 55MPH or lower	One byproduct, a bio-char, can be used in soil to absorb greenhouse gases
1974 savings from lowering speed limit: 0.2–3.0% per year	The pyrolysis process is still experimental and could have drawbacks

Table 2. Five sample “facts” relevant for answering our two task questions (55MPH & Pyrolysis).

Total facts discovered

Some of our analyses included the total number of facts discovered as a dependent variable. Of course, participants uncovered lots of information while searching; here we only tally “facts” from the thirteen we deemed relevant.

Learning scores

Other analyses took into account comprehension of each fact through a 3-point *depth of processing* scale (Table 3), modeled after Bloom’s Taxonomy (Bloom & Krathwohl 1956). Facts that were simply identified or perceived (e.g., subjects reported information from a website) received one point. Facts that were paraphrased or reworded received two points. Facts that were later evaluated in comparison to other information or otherwise synthesized received three points. Ultimately subjects were assigned a single Learning Score (reported later in Table 5): a sum total of the highest *depth of processing* score for each of the 13 relevant facts they discovered. For example, ss02 received a Learning Score of 6 in the Social Condition from receiving *depth of processing* scores of 3 on two important facts.

Score	Definition	Example
1	Identifies or perceives certain facts, data, or information	Reading a fact off the screen.
2	Understands the meaning of information; presents a translation of the information	“So they’re saying that...”; “in other words, this means that...”
3	Integrates and synthesizes learned information	Compares information to other data points; Uses information as part of a final answer.

Table 3. Our 3-point *depth of processing* scale for assessing subjects’ comprehension of each fact discovered.

RESULTS

Our results focus on the process of social exploration and social reasoning that we observed. Since social search is still poorly supported by web technologies, we do not intend to do a full comparison of performance between the Social and NS Conditions. Instead, we focus on the outcomes of various social strategies in the Social Condition.

Performance results

Before turning to these social strategies, we briefly consider the performance differences between tasks and experimental conditions. Our two task questions resulted in different success rates, with Pyrolysis Learning Scores higher during NS than Social Conditions; in contrast, the 55MPH question resulted in approximately equal performance across conditions (Table 4). This suggests that certain queries are better suited for social search than others. The 55MPH question may have elements of common sense that friends are able to reflect on from personal experiences, whereas the Pyrolysis question may be too technical or esoteric for casual colleagues—or even energy and environmental experts—to have specific knowledge about.

Similarly, we saw performance differences between conditions; namely Learning Scores were higher during NS blocks (Mean: 11.1, Std. Dev.: 6.4) than Social blocks (Mean: 6.8, Std. Dev.: 5.2) (Table 4). In many ways, it is not unexpected that traditional web search still has advantages over social searching on its own—after all, search algorithms have been continually enhanced for the past fifteen years. At the same time, it’s also possible that social tactics require a longer time-course for success, going beyond what we would have observed in our experimental sessions.

	Social	Non-Social	Row Means (Std. Dev.)
55MPH	9.0	11.3	9.9 (6.2)
Pyrolysis	3.0	11.0	8.0 (6.2)
Col Means (Std. Dev.)	6.8 (5.2)	11.1 (6.4)	

Table 4. Mean Learning Scores by task question for the Social and Non-Social Conditions. Column Means report average scores for the two conditions; Row Means report average scores for the two task questions.

These results prompted us to consider what factors might correlate with Learning Scores. There were time differences between tasks: participants in the Non-Social Condition spent between 6 to 20 minutes searching (average: 10 minutes); in the Social Condition, they spent between 12 to 36 minutes (average: 19 minutes). Surprisingly time spent searching did not correlate with performance outcomes in either condition. Additionally, none of the pre-test survey responses appeared to influence performance, and subjects’ performance across conditions was not correlated.

Behavioral patterns

Just as individual performance varied, so too did information seeking strategies. Six general categories of behavior emerged from our analysis: *Targeted Asking*, *Network Asking*, *Searching*, *Checking for Replies*, *Thinking*, and *Other*.²

Targeted Asking, *Network Asking*, and *Searching* were employed as primary social tactics in the Social Condition. In *Targeted Asking*, users identified specific friends or colleagues to ask for help one-on-one. This interaction occurred over email (asynchronously) or over instant messaging (IM) or telephone conversations. *Network Asking* involved posting a question in a public venue, typically an organized social networking site (e.g., Twitter, Facebook), although one user [ss03] posted his query on question-answer sites (Ask MetaFilter and Aardvark³).

Of course, *Searching* occurred in both Social and NS Conditions. When users were restricted to only social resources, they were permitted to search over repositories of social data (e.g., archived data on question-answer sites or Twitter). In the NS condition, they readily exploited Google, Yahoo!, and Wikipedia search functionality.

Checking for Replies was common in the Social Condition for users who engaged in either Targeted or Network Asking. One subject was regularly overhead saying: “I’m going to look and see if I’ve gotten any comments” or “Let’s see if I’m getting any response” [ss06]. *Thinking* was observed in talk-aloud segments where users were contemplating an answer or synthesizing information they had discovered. Episodes where subjects were overtly socializing with friends or chatting with the researchers were coded as *Other*. This suite of strategies is reflected in the color coding of behaviors on ss03’s timeline (Figure 1). An elaboration of his process mapped to behaviors on the timeline is presented below.




Figure 1. Timeline of search activities for subject ss03. Color codes match the suite of behaviors that were common across subjects.

At the start of the task, ss03 goes immediately into a social networking site, Ask MetaFilter. This network, while familiar, had not been visited in about a year, but it was considered to be the “go-to” site for question-answering. He begins by searching for previously posted question-answer pairs on MetaFilter [beige coding]. After gathering reasonable background on the domain, he composes his query. Posting the question was actually somewhat involved [see long blue phases]: He reframes and rephrases his query several times before finally posting it.

In the process of asking his question on MetaFilter, ss03 realizes that the question contains sub-parts that could be outsourced to other communities. He then posts one sub-question to Twitter and another to Aardvark. Next, after searching his mental address book and iPhone contact list for friends he can solicit for help, he texts a few candidates with his question, then moves into Google Chat (an IM client). He soon begins IMing with another friend [green bands] whose status indicator had signaled his availability.

Interspersed with these activities, ss03 occasionally checks for replies to his network questions [yellow], chats with the researchers [red], and searches for information (briefly) on Yahoo! Answers [the beige bands towards the end]. By the end of the session, he has accumulated several facts that are relevant to answering the question and concludes by synthesizing everything he has learned into a final answer [gray phases towards the end].

This demonstrates how a suite of tactics might be employed in a single search session. Of course, some people didn’t have a friend to call, or felt uncomfortable asking questions of their friends on Facebook. A few never even thought to search in a social database. As a result, participants’ use and combination of social strategies were highly diverse.

“Go-to” sources

There were some qualitative similarities across approaches: Participants who had “go-to” sources in mind always turned to them first. Often these “go-to” sources were particular people. For example, ss08 “started with David, [her] know-it-all friend,” and without hesitation sent him an email request. Subject ss02 also initially turned to a friend, Sarah, who was involved in “sustainable studies and environmental work,” a domain likely considered to be relevant for our energy policy

² Targeted Asking, Network Asking, and Checking for Replies were not relevant in the Non-Social Condition.

³ Aardvark is an IM-based question-answer service. Questions are routed to knowledgeable people from users’ extended social network.

questions. This speculation is supported by a later comment that he was looking for people who “may actually know the answer to this” [ss02]. This user was so committed to getting a hold of Sarah, he first emailed her but followed up by phone five minutes later.

For others, these “go-to” sources were particular sites or tools, often tapping familiar social networking sites at the outset [ss01, ss03, ss06, ss07]. Subject ss01 commented: “The first thing I would do for this, because I pretty much go to ‘lazyweb’—i.e., Twitter—all the time, is just go to Twitter.” Similarly, ss06 started by sending the question “out to [about 30] of the networks that [he is] a part of” through the service Ping.fm. Since these participants have large and active social networks and routinely solicit them for information, it is likely that these networks have become integral to their social search strategies already as “go-to” sources.

Social relationships and social decisions

Subjects’ relationships with sources also affected their information seeking decisions. Several participants commented that they sought “knowledge mavens” or domain experts in related fields. When such target candidates existed, they were the “go-to” sources initially solicited for help. However, one-on-one interactions occurred throughout the search sessions. When friends were recruited at later points in the process, it was more likely because they were presently available or willing to engage in problem solving. This general finding agrees with Borgatti and Cross (2003) and Cross and Sproull (2004), that information seeking in social contexts is not only about finding the nearest domain expert.

Many factors influence who searchers turn to for help. We are calling these *social decisions* after the set of *social* factors that go into the decision to use certain social resources during search. Our participants’ social decisions matched those laid out in the formal model of social information seeking by Borgatti and Cross (2003). These include perceptions of:

1. a person’s knowledge and authority “*I trust what she says*” [ss06]
2. their accessibility “*I know he’s online and I know he’s at home*” [ss01]
3. the social costs or obligations that would be incurred “*He’s gonna hate me for this...bugging him about stuff that’s not relevant to our usual interaction*” [ss03]

We noted several additional factors, as well: the history of the relationship, time since the last interaction, and the technology used to communicate (in the present and historically). For example, ss08 chose David not only because he was a “know-it-all friend,” but also because: “We have an engagement over email where we can do this kind of thing.” Subject ss01 reached out to Alex because she had “been IMing with [him] recently.” Another [ss07] used her Facebook buddy list to find someone to ask: “I can see who’s online right away...who’s even on here?”

Understanding how these social decisions relate to users’ overall objectives in social search tasks is absolutely critical for developing supporting tools. While our study cannot address the reasons for turning to social resources in the course of natural search tasks, it does begin to illuminate some of the cognitive consequences of social interactions in search.

Outcomes of Social Strategies

We observed several common benefits from seeking information from targeted friends and wider networks.

Informational benefits of social resources

First, friends provided information (facts) to our users. Both Targeted and Network Asking had a relatively high rate of return for 20 to 30 minute sessions. Although responses came at a time delay, subjects received between 1 and 6 replies from friends. In fact, the number of replies received was correlated with subjects’ Learning Scores (Spearman R: 0.79; $p < 0.05$), suggesting that friends can, in fact, play an important role in supplying information.

Initially, we expected that certain social tactics would emerge as superior to the others, as each of the three strategies (Targeted Asking, Network Asking, and Searching) has different merits. Targeting individuals can be useful if they are knowledgeable and available to respond. Querying a network distributes a question over a wide (diverse) audience, theoretically increasing the likelihood of reaching an individual with the appropriate knowledge and availability. Searching, of course, could provide ample information (with large databases), but is limited to the content already present in the database.

However, no individual tactic was superior to others; instead, combining these social tactics consistently led to significant performance gains. Performance outcomes as a function of social tactics are reflected in both the *total number of facts discovered* in the session (Spearman’s R: 0.77; $p < 0.05$) and *depth of processing* of information via Learning Scores (Spearman’s R: 0.73; $p < 0.05$) (Table 5). In other words, people who employed more strategies uncovered and synthesized more task-relevant information. Additionally, the number of social tactics used correlated with task performance better than any other measure we collected, including number of friends on the primary social networking site, social network diversity (as measured by the *position generator*), and background knowledge and interest in energy policy. This suggests that the

potential availability of and access to knowledge within a network may not determine the search outcome—users must know how to manipulate their social environments. Targeted and Network Asking are types of social manipulations, and those who adeptly exploited their available social resources outperformed those who didn't (regardless of other network factors).

ID	#social tactics	#facts discovered	Learning Score
ss04	1	1	1 (Pvto)
ss07	1	1	3 (55)
ss02	1	2	6 (Pvto)
ss06	2	1	2 (Pvto)
ss05	2	5	9 (55)
ss08	3	3	7 (55)
ss01	3	6	9 (55)
ss03	3	8	17 (55)

Table 5. Combining social tactics corresponded to better task performance (# facts discovered and Learning Scores).

At the same time, it is possible that Targeted and Network Asking are manipulations of the same social environment, but in complementary ways (e.g., Network Asking could be used to identify people for subsequent Targeted Asking). These approaches could be partnered in theory, but in practice, most participants used each social tactic as distinct strategies. For example, ss08 first emailed her friend David, then searched over Yahoo! Answers, and finally asked a question on Twitter. Her trajectory through the search space was fluid and natural, but her efforts were chunked into three distinct phases.

Cognitive benefits of social resources

Yet friends can provide more than just facts and information—they can also act as problem-solving aids (Cross & Sproull 2004). Such cognitive benefits differed depending on our subjects' strategies. Searching (over social databases) produced relatively few task-relevant facts and rarely resulted in episodes of contemplation and deep thinking. On the other hand, Targeted and Network Asking led users to become cognitively engaged with task material, but in different *question-answering* phases.

Question composition prompts thinking in Network Asking

When interacting with social networking sites, several subjects realized cognitive benefits during the first phase of social search—the *questioning* phase. Three out of five users who employed Network Asking engaged in episodes of deliberate thinking and contemplation while they were composing their query. Their talk-aloud verbalizations revealed this [ss03]: “Okay, what can I ask as a quick thing that somebody might actually reply to?” He later paused after the first draft of his query (for the 55MPH question):

So that's a big open question. We still don't know how many barrels we use per year; we do know how many gallons are in a barrel, so we need to know...[pause]. The problem is—the number of barrels per year kinda helps, but it's not the full answer, cause I don't know how much of that is...consumed by cars.

Another subject similarly typed out the first draft of a question in Twitter, only to delete it and consider: “Let's see...what do I really want to be asking?” [ss05]. These episodes of thinking around query composition were not prevalent during Targeted Asking, which is somewhat surprising since the communication modes of email and IM offer similar technological affordances (small text box, asynchronous exchanges) but in private channels. Perhaps posting a public and archival question prompts people to more seriously consider the nature of their query? If this is the case, social networking sites may, curiously, provide benefits to the framing and formulation of search questions.

Private replies prompt thinking during Targeted Asking

In contrast, Targeted Asking offered cognitive support primarily in the *answering* phase of social search. For the five subjects who engaged in Targeted asking, all paused to synthesize information after receiving private replies (from both synchronous or asynchronous exchanges). One subject [ss02] commented after the phone call with Sarah:

She started by saying [pyrolysis] is probably a way to convert from one liquid or substance to another. Well, breaking down the word, the first part of the word means fire or inflammable...

...so maybe it's a medium in which...you convert the substance from one thing to another, and it reduces or eliminates the carbon emissions.

Similarly, ss08 evaluated two conflicting answers about fuel economy and vehicle speed (15–25% savings versus no savings):

David is really knowledgeable and he says it's 15–25% savings...But I also believe [Derek], who's an electrical engineer and thinks in terms of mechanics—probably I would average the two. I could assign a higher weight to David's response, in which case I would bear on the higher side of what he suggested, but if I weren't weighing his answer, I would say 15%.

This behavior can also be seen in the color coded activities in subjects' timelines (Figure 2) where the green and gray bands (asking friends and thinking episodes) follow each other. It should be noted that the green–gray pattern is not always present because some of the green blocks mark when users *asked* questions of their friends. Thinking and contemplation were observed after *receiving* private replies.⁴ In some ways, this is expected because social interactions are known to be a source of cognitive development (Rogoff 1990, Pea 1993). Yet this behavior was not observed following public replies on social networking sites, arguably also socially-generated material. Thus, we conclude that private social channels facilitated the search process in complementary ways to public or semi-public networks.

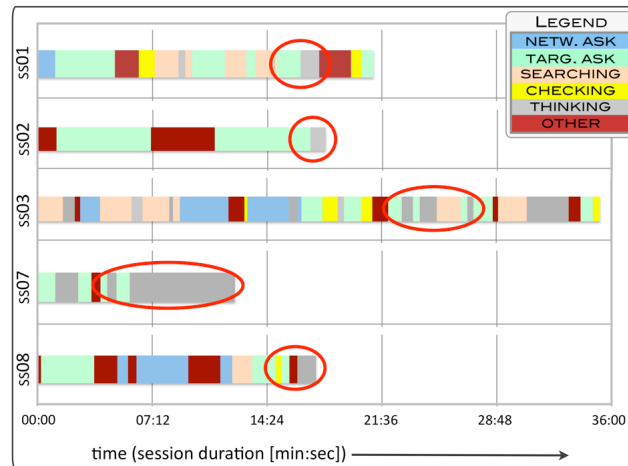


Figure 2. Thinking episodes [gray] follow private replies from friends during targeted asking [green].

Quality of friends' responses

A possible reason for these differences may lie in the quality of the responses received. Replies from social networking sites tended to be *content-light*—predominantly conversational, goofy, or off-target. They did not contain task-relevant facts that helped advance understanding of the search question. For example, one Twitter reply suggested that there would be no gas savings at 55 MPH “because no one drives the speed limit!” Another posited: “Lots more waste sitting idly in traffic.” A third, via Facebook, commented about pyrolytic oil: “Isn't that something I rub on my [body]? Are you still in San Francisco?”

In the days following the study, we contacted the friends who helped our participants in order to assess the motivations behind their responses. There were twelve individuals who provided *public* replies through social networking sites—eight answered our follow-up prompt, reporting that they were mostly looking to start a conversation rather than provide a substantive reply.

I was curious to know why she thought the way she did and more importantly to start a small conversation. [R7]

Seemed rhetorical and random, so I thought I'd reply for fun. [R8]

My reply was more snark than anything else. [R3]

In contrast, replies from targeted friends were *content-rich*. Ss08's friend, David, replied that a 55MPH limit would result in “15–25% savings. But if those cars were electric, we'd have all those barrels left to use for something else.” A friend of ss01 provided details about the complexity of the problem: “There's no one national speed limit, there are two: 55 miles per hour in general, 65 miles per hour for certain roads.”

As before, we followed up with these *private* respondents to learn more about their motives. The three who replied (out of six contacted) reported that they wanted to provide information to their friends; they were less interested in the “social” or conversational aspect of the exchange. Perhaps as a result, our subjects had the incentive and the material for serious reflection.

⁴ This analysis included replies received from targeted friends where we coded for Checking for Replies.

Searching in Social and Non-Social contexts

Because our goal in this study was to understand social exploration and reasoning during search, we devoted the majority of our analysis to the Social Condition. Here we present a brief analysis of the more traditional searching behaviors. Searching was observed in the Social Condition over social databases (e.g., Twitter, Yahoo! Answers, etc.) and as the entire search strategy in the NS Condition.

In many ways, searching across social sources looked very similar to searching on Google and Yahoo!. The search tools were comparable (e.g., single line search box) and global strategies were the same (e.g., submitting our entire task question as a multi-string query). Participants even found many of the same information resources when searching in both conditions.

However, there were qualitative differences in searching outcomes between conditions in the ratio of *search queries* to *thinking episodes*. We considered the number of searches executed as the search query count—searching with Google or Yahoo! [NS Condition] or searching a social database [Social Condition]. Thinking episodes were tallied from our *Thinking* behavioral coding category. From this, we found a higher ratio of *thinking episodes* to *search queries* in the NS than Social Condition. Put another way, searching in social sources resulted in less deliberation per search (mean ratio of thinking/search: 0.29) than using search engines and other non-social tools (mean ratio of thinking/search: 0.53).

These differences can be explained in part by looking at the quantity and quality of the search results. Social sources don't have the scale and scope of traditional search engines, and thus may be limited by providing fewer hits. This was reflected in our data: participants identified fewer task-relevant pages per search when using social databases (mean: 0.67 pages/search) than when using traditional search engines (mean: 1.55 pages/search).

However, pages perceived by the user as relevant resulted in roughly the same amount of thoughtful deliberation regardless of where the material was found (search engines or social tools). The ratio of the number of *thinking episodes* to *relevant pages* was approximately equal across conditions (0.27 in the Social Condition; 0.31 in the NS Condition). This suggests that subjects are able to make accurate judgments on social material, dismissing irrelevant content but paying attention to important facts. If this is the case, the act of searching may not be measurably different in social versus non-social contexts—searchers may simply need to expend more effort finding relevant material in limited social databases.

DISCUSSION

In our study we observed several interesting patterns and noted that different types of social engagements (e.g., public versus private) may provide complementary cognitive benefits. While further study is needed to isolate the exact reasons *why* these benefits occur, we can discuss a few of the possibilities for our findings.

Combining social strategies

Regarding our observation that combining social tactics led to better task performance, we might wonder why all users did not exploit the suite of social options.

One possibility is that early successes become barriers to subsequent discoveries. Several subjects who found (relevant) information early in their process seemed less likely to search in another channel at a later point. When ss02 spoke with Sarah, their conversation was short, succinct, and produced reasonable information about pyrolysis; but the facts were incomplete, and he did not extend his problem solving when the conversation ended. Similarly, ss07 only spoke briefly with a friend on Facebook, but after getting reliable information from him, she ended her information seeking process. (Both received mediocre Learning Scores.) Why might targeting friends for help—and getting good information from them—produce weaker outcomes than exploring multiple social channels? Is one possible shortcoming of “expert” (social) opinions that people place too much trust on their data? Bandura (1989) suggests that “sufficing outcomes can...operate as barriers to discovery.” In other words, “once people hit upon a solution that is sufficient, they keep using it without considering other alternatives, even though better ones [may] exist.” (Schwartz 1982 in Bandura 1989). If such factors affect search as well, system designers could try to increase social presence on a site or reveal alternative social channels to information seekers.

A second possibility is that different social strategies may be optimal for different search phases and information needs. For example, searching in large databases may provide an overview of the problem space (*scope*), but personalized responses may clarify target issues related to the search query (*depth*). Our data didn't reveal specific behavioral patterns to indicate how scope and depth should be paired in a search task. But the nature of exploratory queries suggests that getting the scope and breadth of a domain (i.e., through searching) may be more useful in framing open-ended queries early in the search process, while drilling down to details (i.e., through customized replies) becomes more useful as the problem space becomes more familiar.

It may be noteworthy that none of the social tactics we observed involved *pushing* information to the user: Our subjects' strategies involved *explicitly* seeking information from social communities. This lack of *pushed* data (e.g., recommendations,

tag clouds) is surely an artifact of our experimental design. At the same time, social metadata is often underutilized in search facilities despite a long and rich history with social recommendation and collaborative filtering techniques (Goldberg et al. 1992, Konstan et al. 1997, Glance 2001, Smyth 2007). Google’s “Did you mean” feature and Yahoo!’s Search Assist provide guidance only relative to *common* usage patterns (averages), rarely providing insights into an extended search space. In contrast, MrTaggy (<http://mrtaggy.com>)—a new social search tool—presents *concepts* related to a user’s search query by aggregating keyword tags from social bookmarks, providing scope to users exploring novel or complex domains (Kammerer et al. 2009). Future work could compare how user strategies and outcomes differ when receiving implicit social metadata versus performing active question-answering—Chi’s (2009) distinction between *social feedback* and *social answering*.

Social resources as cognitive aids

It is interesting that Network and Targeted Asking tactics were associated with different cognitive gains during the *questioning* and *answering* phases of search, respectively. Of the many possible reasons for these differences, some may be due to the nature of the tools and the audiences reached.

Network Asking. Social networking sites have a different set of affordances and cultural norms than private messaging channels like email, instant messenger, or—at the other extreme—face-to-face communication. This set of constraints may have led users to reformulate their problem before asking it on a social networking site.

1. One constraint lies in the lack of knowledge about other *participants* in their network. Even on sites like Facebook, with bidirectional, presumably real-life friendships, the audience can grow quite large such that each member’s individual presence is not immediately observable. Perhaps without current knowledge of this audience, our subjects had to carefully craft their information request.
2. Another constraint relates to the *asynchronous* nature of online communication. In synchronous channels (i.e., face-to-face), people can gauge the reaction from listeners to correct blunders, clarify points, and elaborate as necessary. This natural back-and-forth is often lacking in the asynchronous interactions on social networking sites. Instead, there may be increased pressure to “get it right the first time.” Participants may have treated their public question as a one-shot opportunity, not something they could incrementally build upon with collaborators.

Furthermore, questions, comments, and discussions around a topic become archival material in many public forums. The fact that our subjects’ queries were being publicly recorded may have also prompted their early consideration of our task questions.

3. A final constraint comes directly from *cultural norms* instilled in these social networking sites. Many sites encourage user participation, but with short messaging formats. Twitter restricts posts to 140-characters or less. Facebook and Ping.fm don’t have explicit limits, but interface designs and cultural practices likely encourage succinct posts: the “status update” box is (visually) small and users tend to be pithy in their updates and comments. These short messaging practices probably also led our subjects to think about their task problems early on. One subject [ss05] outwardly commented that our question was hard to fit in Twitter’s 140-character limit (“running really close on characters...[can’t] say anything else”); instead, she interpreted and then paraphrased the critical parts of the question, customizing it for a short message format. The process of reworking the question in this way led to insights about what the question was really asking and how to think about it differently.

Targeted Asking. Private messaging channels have an alternative set of affordances that likely contributed to the deeper processing of private replies, as opposed to the deeper formulation of early queries. The near-synchronous, personalized nature of one-on-one responses may underlie these observed cognitive benefits.

1. In contrast to addressing wider networks, targeted messaging usually occurs between people who are (minimally) acquaintances (Cross et al. 2001). This relationship history contributes to a working knowledge of friends’ background and interests. Thus, if the question is asked of the right source, there may be little need to paraphrase or reformulate the initial inquiry; the question will be either familiar or unfamiliar to the friend. As a result, our subjects may have had little motivation to think deeply about their task question in the initial formulation of privately-posed questions.
2. However, the *synchronous* nature of the communication channel may be critical for deep thinking. The majority of our subjects’ targeted interactions occurred via telephone or IM where they received and responded to information in near-to-real-time. Indeed, only after a series of exchanges did users [ss01, ss02, ss03, ss07] achieve the highest *depth of processing* scores on discussed facts. Curiously, these communication channels would be categorized as *moderate social presence media*—remote, synchronous exchanges—by Robert and Dennis (2005). They proposed a paradox of media richness, suggesting that *low social presence channels* (email, fax, voicemail) may actually offer the best opportunities for cognitively processing complex information. We might add that moderate social presence channels may offer similar affordances.

3. Finally, there are different *cultural norms* to observe in private (versus public) communication channels. When asked a specific question, friends may feel obliged to provide a serious answer. In doing so, they likely customize their response to the seeker's perspective (background and knowledge). The content, therefore, becomes much more personalized and relevant than information shared via social networking sites.

Additionally, our subjects may have felt loyalty to friends who were noticeably engaged. If answering a question obviously demanded their time and energy—multiple messages were sent or the exchange lasted a long time—users may be subsequently obliged to think carefully about their replies. Such collaborative involvement may have led to our observations of deeper engagement with privately-shared material.

Limitations

Our results should be qualified by the limitations inherent in our experimental design. First, we collected data from only eight users, each of whom had markedly different interaction styles and task behaviors. Our analysis is based on their 16 search sessions. Second, self-motivated searches are known to have different behavioral profiles than externally-prompted ones (Russell & Grimes 2007, Evans & Chi 2008).

Third, we instructed subjects to focus on the search process rather than try to achieve a certain answer or result. These instructions were intended to elicit a wider variety of social searching behaviors, as well as to allay participants' concerns that their social resources may be limited. Not directing subjects to a specific goal state, however, meant that fine-grained post-test measures were inappropriate; instead we constructed performance measures to capture the exploration process, including depth of processing of search material.

Fourth, our task questions were drawn from only one domain (energy policy), and participants used "general-purpose" social sources to answer them. We may expect to see different social strategies in communities of shared practices or among colleagues with shared interests (e.g., technical forums, domain-specific social networking sites).

Finally, our analysis was restricted to single session search blocks. Exploratory searches—characterized by open-ended queries and frequently requiring multiple iterations—may extend across sessions and involve multiple sources, making them hard to observe in a single sitting. Our strategy for addressing this limitation was to ensure that participants' working environments were as naturalistic as possible, hoping that this would minimally reveal their authentic behaviors and approaches to our task problems. We did not set time restrictions in any session, and allowed participants to signal the natural conclusion of their own search episodes. It is noteworthy that search sessions—in both the Social and Non-Social conditions—were rather long and involved, lasting anywhere from six to thirty-six minutes. Altogether this suggests that we may have captured more than just the early stages of exploratory search with our study's design.

Design Implications

It's unsurprising that different communication channels serve individual cognition in different ways. The design challenge is to create working platforms that integrate the benefits of social communications with those of traditional web search facilities. Our study has begun to illuminate the benefits of engaging social resources during search: Both during the *questioning* and *answering* phases of search, other individuals are useful not only for providing information, but also for triggering new insights to problems. Of course, search engines have an enormous advantage over social strategies on their own because of their vast databases and quick retrieval algorithms. However, social inputs may naturally augment search platforms in cases where the information need is not well served by traditional methods.

Some queries are just better suited for social answering. These may be cases of "vocabulary problems" (Furnas et al. 1987) where users lack the right search terminology (due to mis-framing, having alternative perspectives, or from navigating a novel domain). Similarly, how-to questions often require tacit knowledge from prior experiences; and subjective questions (e.g., dining or lodging recommendations) are addressed through contextual knowledge and personal opinions. Not only may these queries benefit from social inputs, the best results may come from a trustworthy friend over a stranger. A recent report by the PEW Internet & American Life Project (2008) corroborates this sentiment—their results show that community networks are preferred for certain types of information needs (such as choosing a new bank or dealing with a spiritual crisis) (Wells & Rainie 2008).

This presents two challenges for bringing social and collaborative resources to search engines: how to identify the set of queries that benefit from social inputs and how to incorporate social networking and social answering facilities.

First, how do we distinguish the set of queries that will benefit from social inputs? User intent has always been difficult to discern; by classifying queries that are "social" or collaborative in nature, we will begin to have a framework for identifying those use cases in practice. Another approach may be to observe users in action, taking cues from a series of behaviors within a single search session. Users suffering from vocabulary problems may exhibit certain types of behaviors—using long multi-string queries, searching with a "how to" phrase, frequently reformulating search terms in short period of time, or seeking

help from the “Did you mean” (Google) or “Search Assist” (Yahoo!) options. Developing methods for finding patterns in information seeking actions will help in the first stage of social search support: identifying the relevant use cases.

The second challenge is how to effectively integrate social networking, direct messaging, and social recommendations with current search facilities. One suggestion is to begin with traditional search. The great value in search engines is that they can efficiently provide *scope* on a problem domain, quickly revealing whether the answer is buried deep or not. If it becomes apparent that search results do not immediately address the user’s needs (signaled by the behavioral cues above), a search engine could then take explicit actions to suggest alternative approaches. For example, high-level concepts derived from tags and other social metadata could be displayed in a sidebar on the page—as MrTaggy does (Kammerer et al. 2009). This would aid users in navigating a search space by concepts and categories, not just by keywords. Or, available domain experts from the community (or filtered by the user’s address book) could be suggested as social resources: Instant messaging facilities could be conceivably incorporated into search pages to facilitate these directed interactions (i.e., Targeted Asking). Finally, the site could encourage users to tap their social networks, broadcasting their question to a large personal network (i.e., Network Asking).

While our study found cognitive benefits to posing questions to large networks, having one-on-one conversations with peers was better at ensuring quality responses (*social answering*, Chi 2009). An optimal social search program would include both types of facilities. First, many social networking sites have adopted short messaging formats (i.e., status updates), which can serve as a query broadcasting mechanism. Next, expertise-matching algorithms could route questions to certain individuals in the network—either self-identified experts or *experts of convenience* (who are presently available). This combined social strategy may increase the likelihood of a relevant response by (1) reducing noise (well-matched targets will provide well-matched answers) and (2) reducing the bystander effect (Darley & Latané 1968) (where people become less likely to help as more bystanders are present). Although this is a style of asynchronous, remote collaboration, our results suggest that there are informational, and *also cognitive*, benefits to such momentary social collaborations even for users with individual search goals.

CONCLUSION

Our study provides some early insights to the behaviors and practices surrounding question-answering in social environments. Let us briefly summarize:

- We identified three social tactics for information gathering (Targeted Asking, Network Asking, and Searching).
- These tactics used in combination led to better search outcomes (based on Learning Scores).
- Query composition on social networking sites resulted in reformulation and framing of the task problem.
- Processing (targeted) friends’ answers led to greater synthesis of task-relevant material, through episodes of thinking and contemplation.
- Facile manipulation of social environments (through the three social tactics) correlated better with task performance than factors like social network size, network diversity, or background knowledge.
- Yet, Learning Scores were still higher under tradition web searching than social searching alone.

This begs the question: *Do your friends make you smarter?* Our results would suggest *yes*: other people serve as cognitive aids during search and sensemaking tasks. Social resources may be essential for answering a certain class of queries and, in partnership with powerful search engines, could greatly assist users on their quest for information.

The evidence from our study alone is not enough to determine which of the underlying factors are most important for designing integrated social search systems. Future studies should look more closely at question-answering behaviors and outcomes across different social environments. As a research community, we must develop a more sophisticated understanding of how social searches of various types compare to traditional web search. We have several models of social information seeking (Borgatti & Cross 2003, Evans & Chi 2008, Pirolli 2009) and collaborative information seeking (Jentzsch & Prekop 2002, Golovchinsky et al. 2008, Shah 2008) to draw upon. From here, we need direct behavioral observations to understand how these phenomena play out in digital environments: How do personal goals influence social decisions? How do social capital, social presence, and tie strength affect information seeking? Will different online communities be more amenable to question-answering? How do collaborative filtering and social recommendations complement explicit social information seeking practices? We hope that our study provides some initial contributions to a more holistic understanding of question-answering and social search.

ACKNOWLEDGMENTS

We extend a great thanks our willing participants and our collaborators in the Augmented Social Cognition group at PARC. We thank the reviewers of our manuscript for their critical comments. And we are grateful for the guidance from our energy policy experts: Gabriele Cente, Johannes Lehmann, and Lew Fulton. This work was supported in part by Office of Naval Research Contract No. N00014-08-C-0029 to Peter Pirolli.

REFERENCES

1. T.J. Allen. *Managing the flow of technology: Technology transfer and the dissemination of technological information within the R & D organization*. Cambridge, MA: MIT Press, 1997.
2. S. Amershi and M.R. Morris. CoSearch: A system for colocated collaborative web search. In *Proc. CHI'08, ACM Press*, pages 1647–1656, 2008.
3. A. Bandura. *Social cognitive theory*. In R. Vasta (Ed.), *Annals of child development. Vol. 6. Six theories of child development* (pp. 1–60). Greenwich, CT: JAI Press, 1989.
4. B.S. Bloom and D.R. Krathwohl. *Taxonomy of educational objectives: The classification of educational goals, by a committee of college and university examiners. Handbook 1: Cognitive domain*. New York, NY: Longmans, 1956.
5. S.P. Borgatti and R. Cross. A relational view of information seeking and learning in social networks. *Management Science*, 49(4):432–445, 2003.
6. R. Burt. *Structural holes: The social structure of competition*. Cambridge, MA: Harvard University Press, 1992.
7. L. Catledge and J. Pitkow. Characterizing browsing strategies in the World-Wide Web. *Computer Networks and ISDN Systems*, 27(6):1065–1073, 1995.
8. E.H. Chi. Information seeking can be social. *Computer*, 42(3):42–46, 2009.
9. A. Cockburn and S. Jones. 1996. Which way now? Analysing and easing inadequacies in WWW navigation. *International Journal of Human Computer Studies*, 45:105–129, 1996.
10. R. Cross, R.E. Rice, and A. Parker. Information seeking in social context: Structural influences and receipt of information benefits. *IEEE Transactions on Systems, Man and Cybernetics—Part C*, 31(4):438–448, 2001.
11. R. Cross and L. Sproull. More than an answer: Information relationships for actionable knowledge. *Organization Science*, 15(4):446–462, 2004.
12. J.M. Darley and B. Latané. (1968). Bystander intervention in emergencies: Diffusion of responsibility. *Journal of Personality and Social Psychology*, 8:377–383, 1968.
13. D.C. Edelson, R.D. Pea, and L. Gomez. *Constructivism in the collaboratory*. In B.G. Wilson (Ed.), *Constructivist learning environments: Case studies in instructional design* (pp. 151–164). Englewood Cliffs, NJ: Educational Technology Publications, 1996.
14. K. Ehrlich and D. Cash. *Turning information into knowledge: Information finding as a collaborative activity*. In J.L. Schnase, J.L. Leggett, R.K. Furuta and T. Metcalfe (Eds.), *Proceedings of Digital Libraries '94* (pp. 119–125). College Station, TX, 1994.
15. K.A. Ericsson and H.A. Simon. Verbal reports as data. *Psychological Review*, 87(3):215–251, 1980.
16. B.M. Evans and E.H. Chi. Towards a model of understanding social search. In *Proc. CHI'08, ACM Press*, pages 485–494, 2008.
17. E.A. Fox, D. Hix, L.T. Nowell, D.J. Brueni, W.C. Wake, L.S. Heath, and D. Rao. Users, user interfaces, and objects—envision, a digital library. *JASIS*, 44(8):480–491, 1993.
18. G.W. Furnas, T.K. Landauer, L.M. Gomez, and S.T. Dumais. The vocabulary problem in human-system communication. *Commun. ACM*, 30(11):964–971, 1987.
19. H. Ganzeboom, P. de Graaf, and D. Treiman. A standard international socio-economic index of occupational status. *Social Science Research*, 21:1–56, 1992.
20. N.S. Glance. Community search assistant. In *Proc. Intelligent User Interfaces*, pages 91–96, 2001.
21. D. Goldberg, D. Nichols, B.M. Oki, and D. Terry. Using collaborative filtering to weave an information tapestry. *Commun. ACM*, 35(12):61–70, 1992.
22. G. Golovchinsky, J. Pickens, and M. Back. A taxonomy of collaboration in online information seeking. In *JCDL Workshop on Collaborative Information Retrieval*, 2008.

23. M. Granovetter. The strength of weak ties. *American Journal of Sociology*, 78(6):1360–1381, 1973.
24. M. Granovetter. *The strength of weak ties: A network theory revisited*. In P. Marsden & N. Lin (Eds.), *Social Structure and Network Analysis* (pp. 105–130). Beverly Hills, CA: Sage Publishers, 1982.
25. T. Hatch, and H. Gardner. *Finding cognition in the classroom: An expanded view of human intelligence*. In Solomon, G. (Ed.), *Psychological and Educational Considerations* (pp. 164–188). Cambridge University Press, 1997.
26. R. Jentsch, and P. Prekop. *A conceptual model of collaborative information seeking*. In D. White, *Knowledge Mapping and Management* (pp. 98–113), Hershey, PA: IRM Press, 2002.
27. Y. Kammerer, E. Nairn, P. Pirolli, and E.H. Chi. Signpost from the masses: learning effects in an exploratory social tag search browser. In *Proc. CHI '09, ACM press*, pages 625–634, 2009.
28. I. Karasavvidis. Distributed cognition and educational practice. *Journal of Iterative Learning Research*, 13(1/2):11-29, 2002.
29. J.A. Konstan, B.N. Miller, D. Maltz, J.L. Herlocker, L.R. Gordon, and J. Riedl. GroupLens: Applying collaborative filtering to Usenet news. *Commun. ACM*, 40(3):77–87, 1997.
30. J. Lehmann. Bio-energy in the black. *Frontiers in Ecology and the Environment*, 5(7):381–387, 2007.
31. G. Marchionini. Exploratory search: From finding to understanding. *Commun. ACM*, 49(4):41–46, 2006.
32. C. Marlow. The Structural Determinants of Media Contagion. Ph.D. Dissertation, Massachusetts Institute of Technology, School of Architecture and Planning, Program in Media Arts and Sciences, 2005.
33. M.R. Morris. A survey of collaborative web search practices. In *Proc. CHI'08, ACM Press*, pages 1657–1660, 2008.
34. M.R. Morris and E. Horvitz. SearchTogether: An interface for collaborative web search. In *Proc. UIST'07, ACM Press*, pages 3–12, 2007.
35. R.D. Pea. *Practices of distributed intelligence and designs for education*. In G. Salomon (Ed.), *Distributed cognitions: Psychological and educational considerations* (pp. 47–87). New York: Cambridge University Press, 1993.
36. J. Pickens, G. Golovchinsky, C. Shah, P. Qvarfordt, and M. Back. Algorithmic mediation for collaborative exploratory search. In *Proc. SIGIR'08, ACM Press*, pages 315–322, 2008.
37. P. Pirolli. An elementary social information foraging model. In *Proc. CHI '09, ACM Press*, pages 605–614, 2009.
38. L.P. Robert and A.R. Dennis. Paradox of richness: a cognitive model of media choice. *Professional Communication, IEEE Transactions*, 48(1):10-21, 2005.
39. B. Rogoff. *Apprenticeship in thinking: Cognitive development in social context*. Oxford University Press, 1990.
40. D.M. Russell and C. Grimes. Assigned and self-chosen tasks are not the same in web search. In *Proc. HICSS'07, IEEE Computer Society*, page 83, 2007.
41. B. Schwartz. Reinforcement-induced behavioral stereotype: How not to teach people to discover rules. *Journal of Experimental Psychology*, 111:23–59, 1982.
42. C. Shah. Toward collaborative information seeking (CIS). In *Proceedings of 1st Collaborative Information Retrieval Workshop, JCDL*, 2008.
43. J. Shepherd. *The information needs and information seeking behaviour of polytechnic history lecturers*. London: British Library, 1983.
44. B. Smyth. A community-based approach to personalizing web search. *Computer*, 40(8):42–50, 2007.
45. R.S. Taylor, R.S. Question-negotiation and information seeking in libraries. *College & Research Libraries*, 29(3):178–194, 1968.
46. M.B. Twidale, D.M. Nichols, and C.D. Paice. Browsing is a collaborative process. *Information Processing and Management*, 33(6):761–783, 1997.
47. United States Government Accountability Office. Energy Efficiency: Potential Fuel Savings Generated by a National Speed Limit Would be Influenced by Many Other Factors. Washington: Government Printing Office, 2008. (GAO-09-153R).
48. M.P.J. Van der Gaag, T.A.B. Snijders, and H.D. Flap. *Position Generator measures and their relationship to other social capital measures*. In N. Lin. & B. Erickson (Eds.), *Social capital: Advances in research*. Oxford University Press, 2008.
49. A.T. Wells and L. Rainie. The Internet as social ally. *First Monday*, 13(11), 2008. Retrieved from: <http://firstmonday.org/htbin/cgiwrap/bin/ojs/index.php/fm/article/view/2198/2051>.